

On the Relationships between Real Consumption, Income, and Wealth

Michael Palumbo
Federal Reserve Board*

Jeremy Rudd
Federal Reserve Board**

Karl Whelan
Central Bank of Ireland†

August 12, 2002

Abstract

The existence of durable goods implies that the welfare flow from consumption cannot be directly associated with total consumption expenditures. As a result, tests of standard theories of consumption (such as the Permanent Income Hypothesis, or PIH) typically focus on nondurable goods and services. Specifically, these studies generally relate real consumption of nondurable goods and services to measures of real income and wealth, where the latter are deflated by a price index for total consumption expenditures. We demonstrate that this procedure is only valid under the assumption that real consumption of nondurables and services is a constant multiple of aggregate real consumption outlays—an assumption that represents a very poor description of U.S. data. We develop an alternative approach that is based on the observation that the ratio of these series has historically been stable in *nominal* terms, and use this approach to examine two basic predictions of the PIH. We obtain significantly different results relative to the traditional approach.

JEL classification codes: E21, C43

*E-mail: michael.g.palumbo@frb.gov.

**Corresponding author. Mailing address: Mail Stop 80, 20th and C Streets NW, Washington, DC 20551.

E-mail: jeremy.b.rudd@frb.gov.

†E-mail: karl.whelan@centralbank.ie. The views expressed are our own and do not necessarily reflect the views of the Board of Governors, the staff of the Federal Reserve System, or the Central Bank of Ireland.

1 Introduction

A well-known complication that arises when testing theories of consumer spending such as the Permanent Income Hypothesis (PIH) is that these theories usually do not apply to purchases of durable goods. As a result, the typical practice has been to focus instead on expenditures on nondurable goods and services. Of course, because the utility theory underlying models of consumption expenditures involves quantities consumed as opposed to nominal outlays, empirical analyses have generally related measures of *real* spending on nondurables and services to measures of *real* income and wealth. In this paper, we draw attention to an important, but relatively neglected, element in the modeling of these relationships—namely, the procedure used to deflate nominal income and wealth.

The standard approach in the consumption literature—as exemplified in the work of Davidson, Hendry, Srba, and Yeo (1978), Blinder and Deaton (1985), Campbell (1987), Galí (1990), and others—has been to relate real expenditures on nondurable goods and services to the measures of real income and wealth that obtain from deflating the corresponding nominal series with a price index for total consumption expenditures. One can show that this procedure yields a reasonable approximation to the true underlying relationship under the assumption that real consumption of nondurable goods and services is, on average, a constant multiple of total real consumption outlays. In this paper, we document that this standard assumption represents a very poor description of postwar U.S. data. Real outlays on durable goods have consistently grown faster than real purchases of nondurables and services, and the hypothesis that these series share a common trend is strongly rejected. As a result, tests based on this assumption rely on a highly inaccurate approximation to the underlying theoretical relationship.

In light of this problem with the traditional approach, we suggest employing an alternative assumption when using outlays on nondurable goods and services to specify tests of consumption theory. Specifically, we observe that the ratio of *nominal* expenditures on nondurable goods and services to *nominal* expenditures on durable goods has been much more stable over the postwar period; the difference in growth rates of the real series is attributable to a falling relative price for durable goods. We show that, under the alternative assumption of a stable nominal ratio, we can accurately re-formulate the predictions of the permanent income model in terms of the relationship between real expenditures on nondurables and services and real income and wealth, with each of the real series constructed by deflating the relevant nominal series by the deflator for nondurable goods and services

consumption.

To illustrate the importance of deflating income and wealth in a manner that accurately approximates a given theoretical relationship, we show how important features of common empirical implementations of the permanent income model can depend quite sensitively on how deflation is handled. Specifically, we examine two predictions of the PIH. The first concerns the relationship between consumption and financial wealth, a subject of topical importance given the large fluctuations in wealth caused by the stock market movements of recent years. The PIH implies a simple long-run relationship between consumption and financial wealth. We demonstrate that when durable goods are accounted for in our preferred manner, the data reveal a relatively strong and stable relationship between consumption and financial wealth; moreover, the recent behavior of consumption has been approximately consistent with this long-run relationship. In contrast, when the traditional deflation methodology is employed, the relationship between consumption and wealth appears weak, and consumption appears to manifest little response to the run-up in wealth that occurred over the second half of the 1990s.

The second prediction of the PIH that we consider is Campbell's (1987) hypothesis that saving should forecast declines in labor income. In his original paper, Campbell employed the standard assumption that real consumption of nondurables and services is approximately a constant multiple of total real consumption. We demonstrate that, once an updated sample is used, there appears to be little evidence that saving predicts future labor income when the theoretically appropriate deflation procedure is applied; by contrast, using the traditional procedure yields stronger—but apparently spurious—evidence that saving forecasts labor income.

2 Two Predictions of the PIH

In this section, we briefly derive the two predictions of the PIH that we test later. At this point, we follow the standard approach in deriving the hypothesis without paying explicit attention to the separate treatment of durable goods; we return to this issue in Section 3.

2.1 The Budget Constraint

Given our focus on alternative approaches to deflating income and wealth, it is useful to begin with a description of the consumer's budget constraint. Consumers start each period

with a stock of nominal financial wealth, \tilde{A}_t , and the nominal labor income flow, \tilde{y}_t , for that period. (Throughout the paper, we use tildes to denote nominal variables.) These can be used to make purchases of consumption goods, \tilde{c}_t , or invested; assets carried forward receive a nominal rate of return equal to i_{t+1} . The resulting budget constraint can therefore be written as:

$$\tilde{A}_{t+1} = (1 + i_{t+1}) (\tilde{A}_t + \tilde{y}_t - \tilde{c}_t). \quad (1)$$

Consumer utility depends on *quantities* consumed, so macroeconomists tend to re-express the budget constraint in terms of real consumption. To do this, we need to deflate both sides by the aggregate consumption deflator, p_{t+1}^c :

$$\frac{\tilde{A}_{t+1}}{p_{t+1}^c} = \frac{p_t^c (1 + i_{t+1})}{p_{t+1}^c} \cdot \frac{\tilde{A}_t + \tilde{y}_t}{p_t^c} - \frac{p_t^c (1 + i_{t+1})}{p_{t+1}^c} c_t.$$

Defining inflation as

$$\frac{p_{t+1}^c}{p_t^c} = 1 + \pi_{t+1},$$

and defining real wealth, real income, and the real interest rate by

$$A_t = \frac{\tilde{A}_t}{p_t^c}, \quad y_t = \frac{\tilde{y}_t}{p_t^c}, \quad r_{t+1} = \frac{1 + i_{t+1}}{1 + \pi_{t+1}} - 1, \quad (2)$$

yields the following representation of the budget constraint in terms of real variables:

$$A_{t+1} = (1 + r_{t+1}) (A_t + y_t - c_t). \quad (3)$$

2.2 The Permanent Income Hypothesis

The workhorse of modern consumption theory is the rational-expectations version of the PIH, as derived by Robert Hall (1978). This theory assumes the representative consumer wishes to maximize the present value of period-by-period flow subutility $U(c_t)$ (where c_t denotes real aggregate consumption). The first-order condition for this problem is given by

$$U'(c_t) = E_t \left(\frac{1 + r_{t+1}}{1 + \rho} U'(c_{t+1}) \right). \quad (4)$$

Assuming quadratic utility and a discount rate equal to the real interest rate yields Hall's result that consumption follows a martingale process (*i.e.*, that $E_t c_{t+k} = c_t$). To derive the predictions that we test later, we need to combine this Euler equation with the budget constraint. Assuming a constant real interest rate of r , performing repeated substitution

on the period-by-period budget constraint, imposing that $(1+r)^{-t} A_t$ goes to zero in the limit, and taking expectations, it follows that

$$\sum_{k=0}^{\infty} \frac{E_t c_{t+k}}{(1+r)^k} = A_t + \sum_{k=0}^{\infty} \frac{E_t y_{t+k}}{(1+r)^k}.$$

Finally, using the martingale property for consumption yields the following “structural” formulation of the PIH:

$$c_t = \frac{r}{1+r} A_t + \frac{r}{1+r} \sum_{k=0}^{\infty} \frac{E_t y_{t+k}}{(1+r)^k}. \quad (5)$$

Campbell (1987) derived a useful way of illustrating the implications of this equation by defining saving as

$$s_t = \frac{r}{1+r} A_t + y_t - c_t = y_t - \frac{r}{1+r} \sum_{k=0}^{\infty} \frac{E_t y_{t+k}}{(1+r)^k}. \quad (6)$$

In words, saving is defined as labor income y_t plus capital income $\frac{r}{1+r} A_t$, minus consumption. This can be rewritten as

$$s_t = - \sum_{k=1}^{\infty} \frac{E_t \Delta y_{t+k}}{(1+r)^k}. \quad (7)$$

This equation, which Campbell labelled the “saving for a rainy day” formula, tells us that saving should help to forecast future declines in real labor income.

In a related paper, Campbell and Deaton (1989) derived a useful reformulation of the rainy-day formula in terms of the rate of saving out of labor income:

$$\frac{s_t}{y_t} \approx - \sum_{k=1}^{\infty} \beta^k E_t \Delta \log y_{t+k} + \kappa, \quad (8)$$

where κ denotes a constant of linearization. In this expression, the discount factor β equals $\frac{1+\mu}{1+r}$, where μ is the expected log-difference of future values of y_t . Thus, if we assume that y_t is difference-stationary in logarithms, we have that the saving rate is stationary, and that it should be negatively correlated with future growth in real labor income. We will describe tests of this prediction in Section 5.

Finally, the saving rate formulation of the rainy-day equation also implies that

$$\frac{\frac{r}{1+r} A_t + y_t - c_t}{y_t} \sim I(0) \Rightarrow \frac{c_t}{y_t} - \frac{r}{1+r} \frac{A_t}{y_t} \sim I(0). \quad (9)$$

Thus, another prediction of the PIH, implicit in the rainy-day equation, is that in the long run there should be a stable relationship between the consumption to labor income ratio and the financial wealth to labor income ratio. We will examine tests of this prediction in Section 4.

3 Accounting for Durable Goods

A problem that arises when attempting to bring the preceding derivation of the PIH to the data is that the model does not account for the special role played by durable goods. Specifically, the Euler equation (4) does not reflect the fact that consumers derive utility from the services provided by the durable goods they own (as opposed to the act of purchasing the items), and that these services are obtained in periods following the purchase. In this section, we look at two alternative methods that allow researchers to test the PIH when durable goods are present.

Empirical research on consumption usually starts with the assumption that purchases of durable goods are explained by a different theoretical model, and that Euler equations for consumption outlays—such as the Hall martingale prediction—only hold for nondurables and services. However, for many of the predictions of consumption theory, such as the two implications of the PIH derived above, one needs to combine the Euler equation with the intertemporal budget constraint. And while there may be a theoretical case for omitting expenditures on durable goods from the Euler equation, we certainly cannot omit them from the budget constraint because they assuredly constitute a use of funds.

The standard approach to this problem has been to invoke a long-run relationship between total consumption c_t and expenditures on nondurables and services c_t^{nd} , and then to use this relationship to arrive at an approximation to the budget constraint that does not explicitly include spending on durable goods. In this way, one can derive predictions for the dynamics of consumption of nondurables and services that still incorporate the long-run restrictions imposed by the intertemporal budget constraint. The particular long-run relationship between c_t and c_t^{nd} assumed in the papers cited in the introduction (and in many others) is that total real consumption is, on average, a constant multiple of real consumption of nondurables and services. This can be written as

$$c_t = \theta c_t^{nd} + \epsilon_t, \tag{10}$$

where ϵ_t denotes a zero-mean approximation error. This implies a new budget constraint of the form

$$\sum_{k=0}^{\infty} \frac{\theta c_{t+k}^{nd} + \epsilon_{t+k}}{(1+r)^k} = A_t + \sum_{k=0}^{\infty} \frac{y_{t+k}}{(1+r)^k}. \tag{11}$$

So, under the assumption that consumption of nondurables and services follows a martin-

gale, we obtain:

$$c_t^{nd} = \frac{r}{1+r} \left(\frac{A_t}{\theta} \right) + \frac{r}{1+r} \sum_{k=0}^{\infty} \frac{1}{(1+r)^k} E_t \left(\frac{y_{t+k}}{\theta} \right) - \frac{r}{1+r} \sum_{k=0}^{\infty} \frac{E_t \epsilon_{t+k} / \theta}{(1+r)^k}. \quad (12)$$

We see, therefore, that making the standard assumption of a stable long-run ratio of real nondurables and services consumption to total real consumption allows us to apply equation (5)'s structural formulation of the PIH to consumption of nondurables and services with only two minor adjustments: First, we need to use “scaled” versions of real labor income and real financial wealth; and second, there is a potentially autocorrelated (and heteroskedastic) approximation error.¹ In a similar fashion, the two predictions of the PIH discussed in the previous section—the rainy-day feature of equation (8), and the consumption-wealth relationship of equation (9)—can also be applied to c_t^{nd} and measures of real labor income and financial wealth. These real measures can be defined as in equation (2) by deflating the corresponding nominal measures by a price index for total consumption, so long as the income and wealth series are appropriately scaled.

These considerations show that the legitimacy of the standard procedure hinges on whether equation (10) provides a reasonable empirical approximation. As is evident from Figure 1, however, the proposition that c_t^{nd} represents a stable fraction of overall real consumption c_t is clearly rejected in postwar U.S. data: Since the 1960s, this ratio has manifested a significant downward trend.² An important implication for the PIH is that the errors in the re-formulated structural solution, equation (12), will be large and trending, as will the errors in the re-formulated versions of our two test equations, (8) and (9). Hence, standard formulations of the PIH may provide a poor approximation to the data even if the underlying theory is correct.

Figure 1 also allows us to identify the source of the long-term decline in the ratio of c_t^{nd} to c_t . The figure shows that the ratio of *nominal* expenditures on nondurables and services, \tilde{c}_t^{nd} , to total nominal consumption expenditures, \tilde{c}_t , has been relatively stable over the post-war period. In particular, while the nominal ratio appears to exhibit a very slight upward trend, this trend is far less pronounced than the downward trend in the real ratio. Thus, the significant disparity between the long-run growth rates of real expenditures

¹The well-known formula of Hansen and Sargent (1980) implies that if ϵ has an $AR(n)$ representation, then this error term—which is a present discounted value of expected future ϵ values—will have an $MA(n-1)$ representation.

²See part II of the Appendix for a complete description of the data used in this paper.

on durable goods and real consumption of nondurables and services stems almost entirely from a large and ongoing decline in the relative price of durable goods. This trend is part of a broader pattern in which prices for producers' durable equipment have also been declining relative to prices for other goods and services, and is the result of productivity growth for the durable goods sector significantly outpacing that for the rest of the economy.³

Figure 1 therefore indicates that an approximation like

$$\tilde{c}_t = \gamma \tilde{c}_t^{nd} + \tilde{\eta}_t \quad (13)$$

(where $\tilde{\eta}_t$ is a zero-mean—and probably heteroskedastic—approximation error), will provide a far more accurate description of the data than equation (10). Indeed, over the sample period shown in the figure, the average absolute value of the approximation error for the relationship that assumes a stable real ratio (equation 10) is more than four times larger than the average absolute approximation error that obtains if a stable nominal ratio is invoked instead (as in equation 13). This suggests that we should seek to reformulate the PIH under the alternative hypothesis of a stable long-run relationship between nominal expenditures on nondurable goods and services and total nominal consumption outlays. It turns out that this is straightforward to do.

Starting with equation (13), we can rewrite the nominal budget constraint, equation (1), as follows:

$$\tilde{A}_{t+1} = (1 + i_{t+1}) \left[\tilde{A}_t + \tilde{y}_t - \gamma \tilde{c}_t^{nd} - \tilde{\eta}_t \right]. \quad (14)$$

This equation can be re-expressed in terms of real consumption of nondurables and services by dividing through with the deflator for this series, p_{t+1}^{nd} :

$$A_{t+1}^{nd} = (1 + r_{t+1}^{nd}) \left[A_t^{nd} + y_t^{nd} - \gamma c_t^{nd} - \eta_t \right]. \quad (15)$$

In this form, the budget constraint involves real labor income, real financial wealth, and real interest rates, where all real variables are defined relative to the deflator for nondurables and services:

$$A_t^{nd} = \frac{\tilde{A}_t}{p_t^{nd}}, \quad y_t^{nd} = \frac{\tilde{y}_t}{p_t^{nd}}, \quad r_{t+1}^{nd} = \frac{1 + i_{t+1}}{1 + \pi_{t+1}^{nd}} - 1. \quad (16)$$

Finally, the structural solution for the PIH can be written as

$$c_t^{nd} = \frac{r^{nd}}{1 + r^{nd}} \left(\frac{A_t^{nd}}{\gamma} \right) + \frac{r^{nd}}{1 + r^{nd}} \sum_{k=0}^{\infty} \frac{1}{(1 + r^{nd})^k} E_t \left(\frac{y_{t+k}^{nd}}{\gamma} \right) - \frac{r^{nd}}{1 + r^{nd}} \sum_{k=0}^{\infty} \frac{E_t \eta_{t+k} / \gamma}{(1 + r^{nd})^k}, \quad (17)$$

³See Whelan (2001) for a more detailed discussion of the relative price shifts associated with the pattern of sectoral differences in productivity growth.

and so the predictions of the PIH can be reformulated in terms of these new measures of real labor income, financial wealth, and interest rates. Importantly, because this derivation is based on equation (13)—which is empirically much more accurate—the approximation errors in the consumption relationship will be much smaller, and hence do not present a major obstacle to testing the theory.

4 The Consumption-Wealth Relationship

We have demonstrated that the methodology traditionally used in constructing structural tests of the PIH introduces significant measurement error in practice, and have proposed a simple alternative approach under which this error will be smaller. We now explore how inferences about consumption behavior are affected by employing this alternative methodology.

In our first example, we consider the PIH’s prediction of a long-run relationship between the ratio of consumption to labor income and the ratio of financial wealth to labor income. According to the version of the PIH developed above, we have that

$$\frac{c_t}{y_t} - \frac{r}{1+r} \frac{A_t}{y_t} \sim I(0),$$

where the errors from a regression of c_t/y_t on A_t/y_t (and a constant) will reflect expectations about future labor income growth.

We estimated this relationship using consumption of nondurables and services and scaled measures of income and wealth derived under each of our two deflation methods. The first regression we consider is based on the traditional deflation method and takes the form

$$\theta \frac{c_t^{nd}}{y_t} = \alpha_0 + \alpha_1 \left(\frac{A_t}{y_t} \right) + \varepsilon_t, \quad (18)$$

where θ is the “real ratio” scale factor of equation (10) and income and wealth have been deflated by the total consumption price index (as in equation 2). Note that this equation can be rewritten in terms of nominal series as:

$$\theta \frac{\tilde{c}_t^{nd}}{\tilde{y}_t} \frac{p_t^c}{p_t^{nd}} = \alpha_0 + \alpha_1 \left(\frac{\tilde{A}_t}{\tilde{y}_t} \right) + \varepsilon_t. \quad (19)$$

The second regression is based on our preferred method. It takes the form

$$\gamma \frac{c_t^{nd}}{y_t^{nd}} = \alpha_0 + \alpha_1 \left(\frac{A_t^{nd}}{y_t^{nd}} \right) + \varepsilon_t, \quad (20)$$

where γ is the “nominal” scale factor in equation (13) and income and wealth have been deflated by the price index for nondurables and services (as in equation 16). This equation can also be rewritten in terms of nominal variables, as follows:

$$\gamma \frac{\tilde{c}_t^{nd}}{\tilde{y}_t} = \alpha_0 + \alpha_1 \left(\frac{\tilde{A}_t}{\tilde{y}_t} \right) + \varepsilon_t. \quad (21)$$

Note that the variable on the right-hand side is the same in both of these regressions (that is, in equations 19 and 21). Because both methodologies deflate wealth with the same price index that is used to deflate income, the ratio of the two “real” wealth and income series in each case is identical to the ratio of the two nominal series. The difference between the two regressions is in the definition of the dependent variables: In addition to the difference in scaling factors, our preferred approach of deflating all series with the same price index implies that the dependent variable in equation (21) is a multiple of the ratio of *nominal* consumption to *nominal* labor income, while the traditional approach yields a dependent variable that equals the nominal ratio multiplied by a downward-trending relative price term (see equation 19).

Fitted values from these regressions are plotted against the actual data in Figure 2: The upper panel plots the results based on the traditional deflation technique, while the lower panel gives the values obtained when our preferred deflation procedure is applied.⁴ Two noteworthy conclusions emerge from a comparison of these charts.

1. Interpretation of the Consumption-Wealth Relationship: Figure 2 indicates that the estimated consumption-wealth relationship fits much less well when the traditional deflation procedure is used. The regression based on our deflation method has an \bar{R}^2 of 0.711, compared with an \bar{R}^2 of 0.363 under the traditional deflation procedure. Of course, this poor fit would not necessarily cause a researcher to *reject* the permanent income hypothesis. However, it may yield a somewhat misleading picture of what must be going on if the model is, in fact, correct.

For instance, we could attempt to reconcile the PIH with the poor fit shown in the upper panel by positing a specific pattern of expectations concerning future labor income growth over this period; as noted, the theory predicts that these expectations determine the residual in this regression. Alternatively, we might view these results as providing evidence that the relationship between consumption and wealth is unstable, perhaps because

⁴The sample period for these regressions extends from 1954:Q1 to 2000:Q4.

of periodic revisions to r , the perceived average return on assets. Hence, we could explain the experience of the second half of the 1990s (during which actual consumption was significantly below its predicted level) by positing that agents were expecting a substantial slowing of labor income growth, lower future values of real asset returns, or both.

However, our analysis suggests a simpler explanation. Specifically, if this long-run implication of the PIH were approximately correct, then one should *expect* to see exactly the pattern of errors evident in the upper panel, in which the fitted values are systematically too high in the early portion of the sample and too low in the latter portion. This is because the traditional deflation procedure introduces an approximation error into this regression that trends downward over time. Interestingly, under our preferred deflation method (which yields a smaller and more stable approximation error) it appears that consumption, income, and wealth moved roughly in line with their usual relationship during the latter part of the 1990s. Apparently, therefore, it is not necessary to appeal to expectations of slower labor income growth or lower expected asset returns in order to explain the behavior of consumption over this period.

2. Estimates of Wealth Effects: Regressions of this type have often been employed to obtain empirical estimates of the long-run marginal propensity to consume out of wealth, which in turn have been used in applied policy analysis. It turns out that these estimates are sensitive to the deflation method we choose: When we fit equation (21) using our preferred deflation methodology, we obtain an estimate of α_1 equal to 0.070, compared to a point estimate of 0.049 under the traditional deflation method.

In addition, estimates of the propensity to consume out of wealth using our preferred method turn out to be relatively stable across different samples, while estimates based on the traditional method are very sensitive to the choice of sample. Inspection of the lower panel of Figure 2 shows that the predicted consumption-income ratio from the regression using our deflation procedure is systematically too high up to 1975, and then generally too low afterward. However, re-estimating this equation over each subsample gives estimates of the propensity to consume out of wealth of 0.082 for the early period, and 0.068 for the later period—figures that are not too far apart. In contrast, estimates obtained under the traditional deflation technique imply propensities to consume out of wealth that are very different over these subsamples—0.124 over the early period, and 0.032 over the later period. Once again, this apparent instability may reflect the presence of a trending approximation

error, which drives a wedge between the true underlying relationship and its empirical counterpart.

5 Saving for a Rainy Day

The preceding results demonstrate that the approximation errors induced by the traditional approach to deflating income and wealth can seriously affect the interpretation of one type of relationship implied by the permanent income hypothesis. Left unanswered, of course, is whether the hypothesis itself is valid. We now consider how tests of a standard prediction of the PIH are affected by the choice of deflation procedure. Specifically, we examine the prediction that saving should Granger cause future declines in labor income, where saving is defined as in equation (6):

$$s_t = \frac{r}{1+r}A_t + y_t - c_t.$$

In addition to examining the effect of the two deflation methods, we will also consider two different approaches to testing this hypothesis that are based on different ways of measuring $\frac{r}{1+r}A_t$ (which corresponds to the real capital income received by consumers).

5.1 Tests Using NIPA Capital Income

Campbell (1987) and Campbell and Deaton (1989) both examined the question of whether saving Granger causes labor income growth. In both papers, $\frac{r}{1+r}A_t$ was equated with personal capital income as measured in the National Income and Product Accounts (NIPAs); hence, the disposable income series employed was identical to the NIPA measure. The specific test we consider here follows the formulation of the Campbell-Deaton paper, which examined whether lags of the saving *rate* (defined as the ratio of saving to labor income) help to forecast the log-difference of real labor income in equations like

$$\Delta \log y_t = A(L) \Delta \log y_{t-1} + B(L) \left(\frac{s_{t-1}}{y_{t-1}} \right), \quad (22)$$

where $A(L)$ and $B(L)$ denote polynomials in the lag operator.

Under the traditional deflation procedure, the Campbell saving rate is measured as

$$\frac{s_t}{y_t} \equiv \frac{y_t^d - \theta c_t^{nd}}{y_t}, \quad (23)$$

where y_t^d is real NIPA *disposable* income and y_t is real labor income, both defined relative to the total consumption deflator. Again, θ is the “real” scaling factor from equation (10),

and c_t^{nd} is real nondurable and services consumption. We use these definitions of the saving rate and real labor income to fit versions of equation (22) that include various lags of each variable; estimation runs from 1954:Q1 to 2000:Q4.

Panel A of Table 1 presents the estimated coefficients on lagged labor income growth and the lagged saving rate (along with the p -values from tests of the exclusion restrictions) that we obtain from this exercise. Starting with the results in the first column (which are from the specifications with one lag of labor income growth and one lag of the saving rate), we see that the saving rate enters the regression for labor income growth with the negative coefficient that the theory predicts; moreover, the coefficient is statistically significant at close to the one percent level. Similar results obtain when we allow for four lags of labor income growth (the second column) as well as when four lags of labor income growth and four lags of the saving rate are both included in the regression (the third column); in each case, the coefficient (or sum of coefficients) on the saving rate is negative and highly statistically significant. Hence, the results obtained using the traditional deflation method confirm the original findings of Campbell and Deaton (1989) that higher saving appears to anticipate slower labor income growth.

This conclusion changes significantly when we employ our preferred deflation method. Here, the saving rate is defined as

$$\frac{s_t}{y_t} \equiv \frac{\tilde{y}_t^d - \gamma \tilde{c}_t^{nd}}{\tilde{y}_t}, \quad (24)$$

where \tilde{y}_t^d is nominal NIPA disposable income, \tilde{c}_t^{nd} is nominal consumption of nondurables and services, \tilde{y}_t is nominal labor income, and γ is again the “nominal” scale factor in equation (13). (Note that this is equivalent to deflating all nominal variables by the price index for nondurable goods and services consumption.) Results from these models are summarized in panel B of the table. For the regressions with one lag of the saving rate (reported in columns one and two), we find that the point estimates of the coefficient on the saving rate, while negative, are smaller and no longer statistically significant. As a result, these specifications provide much less evidence that saving Granger causes income growth. When we allow four lags of each variable to enter the regression (column three), we find somewhat stronger evidence of Granger causality in the sense that the coefficients on the four lags of the saving rate are jointly significant at roughly the one percent level. However, the coefficients on the saving rate essentially sum to zero (and this sum is statistically insignificant), implying that a higher level of saving does not tend to signal lower labor

income growth in the future. It is difficult to argue that this result provides support for the PIH; in any case, these results stand in sharp contrast to those obtained using the traditional deflation procedure.

Another interesting result (which is not reported in the table) is that *both* deflation procedures suggest that increase rates of saving anticipate lower labor income growth when the forecasting regressions are estimated up to 1985 (the end of the sample considered by Campbell and Deaton). Apparently, therefore, it is the introduction of the late 1980s and 1990s that weakens the evidence for Granger causality. Figure 3, which plots the saving rates obtained under the alternative and traditional deflation procedures, helps to explain why.

After the mid-1980s, the saving rate generated by our alternative procedure falls to very low levels, mirroring, for example, the substantial decline in the official NIPA saving rate (which is also calculated as a ratio of nominal series). In contrast, the saving rate calculated using the traditional deflation method declines by a much smaller amount. Labor income growth did in fact pick up in the 1990s, but not by nearly as much as would be implied by the saving measure derived under our deflation procedure. Hence, the traditional method's measure of the saving rate appears to do a better job in forecasting real labor income.

Of course, this saving rate remains quite high during the 1990s for an unsatisfactory reason—namely that, relative to the theoretically correct measure of saving that would forecast labor income growth if the PIH were correct, this series contains an upward-trending approximation error associated with the traditional treatment of durable goods. Put differently, the fact that nominal nondurables and services consumption has remained a roughly stable share of total nominal consumption implies that declining prices of durable goods have resulted in a proportional increase in real durables consumption. When we apply the traditional deflation technique, these declines in durable goods prices boost real income; however, this additional real income is not measured as being spent on real consumption (which is defined as scaled-up real nondurables and services expenditures). As a result, under the traditional deflation method, the measured saving rate is propped up over the 1990s.

5.2 Tests Using Regression-Based Measures of Capital Income

One possible explanation for our failure to find evidence of Granger causality from saving to labor income growth could be that the NIPA capital income series provides a poor proxy for

the appropriate theoretical concept of real capital income, $\frac{r}{1+r}A_t$. We note several possible reasons for this.

First, the correct theoretical rate of return on wealth should include capital gains. However, the national accounts measure of capital income includes only flow payments such as dividends and interest. Thus, a capital-gains driven increase in wealth that is unaccompanied by higher dividends—such that observed in the 1990s—would not enter measured capital income.⁵

Second, even if the current-period flow of *nominal* capital income in the NIPAs corresponded exactly to $\frac{i}{1+i}\tilde{A}_t$ (where i denotes the nominal interest rate), we will not be able to recover the desired real capital income concept $\frac{r}{1+r}A_t$ by deflating the NIPA measure by a price index (as Campbell and Deaton do). Rather, we must subtract a term equal to $\frac{\pi}{1+i}A_t$ —where π is the average (anticipated) inflation rate—from deflated nominal capital income as well.

Finally—and perhaps most importantly—the formulation of the rainy-day equation that we have been working with has been derived from the assumption that nondurables and services consumption follows a martingale. But this result relies on very specific assumptions about preferences; namely, that preferences over nondurables and services are strongly separable from those over durable goods, that the flow utility obtained from nondurables and services consumption is quadratic, and that the real interest rate is constant and equal to the rate of time preference. These assumptions are rather restrictive; more importantly, they are inconsistent with the empirical facts that consumption grows over time and also that the ratio of nominal expenditures on durable goods to nominal outlays on nondurables and services is roughly constant.

It turns out, however, that we can capture the observed long-run stability of nominal expenditure shares if we make the assumption that consumers' preferences are log-linear; *i.e.*, that consumers maximize a lifetime utility function of the form

$$E_t \left[\sum_{k=0}^{\infty} \left(\frac{1}{1+\rho} \right)^k (\alpha_1 \log D_{t+k} + \alpha_2 \log c_{t+k}^{nd}) \right], \quad (25)$$

where D_t denotes the consumer's stock of durable goods. In an appendix, we demonstrate

⁵In his 1987 paper, Campbell associated the real return r as exclusive of unanticipated capital gains (see Campbell, 1987, footnote 9). However, *anticipated* capital gains are also excluded from NIPA measures of disposable income.

that these preferences permit us to derive the following approximate relationship:

$$\frac{\rho}{1+r^{nd}} \frac{A_t^{nd}}{y_t^{nd}} - \gamma \frac{c_t^{nd}}{y_t^{nd}} \approx - \sum_{k=1}^{\infty} \beta^k E_t \Delta \log y_{t+k}^{nd} + k. \quad (26)$$

In other words, once we allow for log-linear preferences and (possibly) growing consumption, we still obtain a long-run linear relationship between the scaled consumption-income ratio and the wealth-income ratio, with deviations from this relationship being related to expectations about future labor income growth. However, in this case the constant multiplying wealth is $\frac{\rho}{1+r^{nd}}$ rather than $\frac{r}{1+r}$.

These considerations imply that, strictly speaking, we should not expect standard saving measures to forecast labor income even if the PIH is true. However, a prediction that is common to both versions of the PIH that we have derived is that the errors from a regression of the scaled consumption-income ratio on the wealth-income ratio should forecast labor income.⁶ This regression-based measure of the “saving rate” can in turn be used in the forecasting regression (equation 22). The regression-based measure is plotted in Figure 4, together with the NIPA-based saving rate that we obtained earlier under our preferred deflation procedure. As expected, the higher capital income that is imputed from the regression-based approach keeps this measure of the saving rate from declining sharply during the 1990s; indeed, the regression-based measure actually *rises* over the latter half of the decade.

We examined whether this variant of the saving rate is better able to forecast labor income growth by using it to repeat the analysis of Section 5.1; the relevant results are summarized in Table 2. Perhaps surprisingly, we are once again unable to find strong evidence that the saving rate Granger causes labor income growth: As the table indicates, we fail to reject the hypothesis that the regression-based saving rate can be excluded from the regressions for labor income growth at conventional significance levels. In addition, the sign of the coefficient (or sum of coefficients) on the saving rate is positive, not negative as the theory would predict. Evidently, then, some additional extension is needed in order to salvage this prediction of the PIH.⁷

⁶Specifically, the negative values of these residuals correspond to the left-hand side of equation (26).

⁷One candidate might be the explicit modelling of anticipated time-variation in asset returns, which we have not allowed for.

6 Conclusions

This paper has examined the well-known problem of how to test theories of consumer behavior when consumption expenditures include durable goods purchases. In particular, we have presented theoretical and empirical arguments for relating real consumption of non-durable goods and services to measures of real income and wealth that are defined relative to a price index for nondurables and services consumption; this contrasts with the usual procedure of deflating income and wealth with a price index for total consumption. In two empirical exercises, we demonstrated that this choice of deflation method can significantly affect the interpretation of observed consumption behavior as well as the results obtained from standard tests of the predictions of the permanent income hypothesis.

Beyond the substantive results relating to tests of consumer behavior, a more general lesson that we take from these findings is that macroeconomists may need to be somewhat more careful regarding their treatment of “real” variables. It is perhaps understandable that economists, who are generally schooled in the dictum that real variables “control for increases in the price level,” might conclude that deflation by a broad-based price index is always the appropriate way to construct a real income, output, or wealth series. However, our analysis shows that this practice can sometimes result in a poor empirical approximation to the underlying theoretical relationship that we seek to capture.

References

- [1] Blinder, Alan and Angus Deaton (1985). "The Time Series Consumption Function Revisited." *Brookings Papers on Economic Activity*, 2, 465-521.
- [2] Campbell, John (1987). "Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis." *Econometrica*, 55, 1249-73.
- [3] Campbell, John and Angus Deaton (1989). "Why is Consumption So Smooth?" *Review of Economic Studies*, 56, 357-374.
- [4] Davidson, James E. H., David F. Hendry, Frank Srba, and Stephen Yeo (1978). "Econometric Modelling of the Aggregate Time-Series Relationship between Consumers' Expenditure and Income in the United Kingdom." *Economic Journal*, 88, 661-692.
- [5] Galí, Jordi (1990). "Finite Horizons, Life-Cycle Savings, and Time-Series Evidence on Consumption." *Journal of Monetary Economics*, 26, 433-452.
- [6] Hall, Robert E. (1978). "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*, 86, 971-987.
- [7] Hansen, Lars Peter and Thomas J. Sargent (1980). "Formulating and Estimating Dynamic Linear Rational Expectations Models," *Journal of Economic Dynamics and Control*, 2, 7-46.
- [8] Whelan, Karl (2001). "A Two-Sector Approach to Modeling U.S. NIPA Data." *Journal of Money, Credit, and Banking*, forthcoming.

Appendix

This appendix is divided into two parts. The first contains a technical note that derives the PIH model under log-linear preferences, as described above in Section 5.2. The second part is a data appendix containing the sources and definitions of the series used in the empirical analysis of Sections 4 and 5.

I. The PIH with Log-Linear Preferences

In this portion of the appendix we derive a microfounded version of the permanent income model in which nominal nondurables and services consumption accounts for a constant fraction of total nominal consumption expenditures in steady state *and* we obtain a closed-form expression for nondurables and services consumption that is similar to equation (17).

The consumer maximizes a lifetime utility function of the form

$$E_t \left[\sum_{k=0}^{\infty} \left(\frac{1}{1+\rho} \right)^k (\alpha_1 \log D_{t+k} + \alpha_2 \log c_{t+k}^{nd}) \right],$$

where D_t is the consumer's stock of durable goods at time t and c_t^{nd} denotes real consumption of nondurable goods and services. Maximization proceeds subject to two constraints. First, we have an equation relating the stock of durable goods to flow real purchases of durable goods, c_t^d , and the rate δ at which durable goods depreciate,

$$D_t = c_t^d + (1 - \delta)D_{t-1}.$$

Second, the evolution of nominal assets is governed by

$$\tilde{A}_{t+1} = (1 + i_{t+1}) (\tilde{A}_t + \tilde{y}_t - p_t^{nd} c_t^{nd} - p_t^d c_t^d).$$

In this latter expression, i_t is the nominal interest rate, p_t^{nd} is the price of nondurable goods and services, p_t^d is the price of durable goods for sale in period t , and tildes denote nominal variables.

As before, the asset accumulation constraint can be written in real terms:

$$A_{t+1} = (1 + r_{t+1}^{nd}) (A_t^{nd} + y_t^{nd} - c_t^{nd} - p_t c_t^d),$$

where p_t is the *relative* price of c_t^d and c_t^{nd} (that is, $p_t \equiv p_t^d/p_t^{nd}$), and we have once again deflated real variables with the price index for nondurables and services. If we define the

variable N_t as $A_t^{nd} + y_t^{nd} - c_t^{nd} - p_t c_t^d$, then we can write the value function for the problem as

$$V_t(A_t^{nd}, D_{t-1}) = \max_{N_t, c_t^d} \left\{ \alpha_1 \log(c_t^d + (1 - \delta)D_{t-1}) + \alpha_2 \log(A_t^{nd} + y_t^{nd} - N_t - p_t c_t^d) + \left(\frac{1}{1 + \rho} \right) E_t V_{t+1} \left((1 + r_t^{nd})N_t, c_t^d + (1 - \delta)D_{t-1} \right) \right\}.$$

The first-order condition for N_t , combined with the envelope condition for A_t , implies the following Euler equation for nondurable goods and services consumption:

$$\frac{1}{c_t^{nd}} = \left(\frac{1}{1 + \rho} \right) E_t \left(\frac{1 + r_{t+1}^{nd}}{c_{t+1}^{nd}} \right).$$

To derive a modified structural version of the permanent-income model, we assume that the expected real interest rate takes a constant value of r^{nd} . It can then be shown that log consumption of nondurables and services approximately evolves according to

$$\Delta \log c_{t+1}^{nd} \approx (r^{nd} - \rho) + \varepsilon_{t+1},$$

with ε_{t+1} denoting a mean-zero expectational error. In addition, the assumption that period-by-period flow subutility takes the Cobb-Douglas form ensures that nominal purchases of nondurable goods and services will be a constant share of total nominal expenditures along a steady-state growth path.⁸

Noting that $r^{nd} - \rho$ is the expected average growth rate of consumption of nondurables and services, we obtain that $E_t c_{t+k}^{nd} = (1 + r^{nd} - \rho)^k c_t^{nd}$. This yields the following analogue to equation (17):

$$c_t^{nd} = \frac{\rho}{1 + r^{nd}} \left(\frac{A_t^{nd}}{\gamma} \right) + \frac{\rho}{1 + r^{nd}} \sum_{k=0}^{\infty} \frac{1}{(1 + r^{nd})^k} E_t \left(\frac{y_{t+k}^{nd}}{\gamma} \right), \quad (27)$$

where once again γ is the factor used to scale nominal nondurable goods and services consumption into aggregate consumption expenditures (we have ignored the presence of the term in the approximation error η_{t+k}).

If we define saving as

$$s_t \equiv \frac{\rho}{1 + r^{nd}} A_t^{nd} + y_t^{nd} - \gamma c_t^{nd}, \quad (28)$$

⁸See Whelan (2001) for a derivation of this result in the context of a multisector general-equilibrium growth model.

we obtain that

$$y_t^{nd} - s_t = \frac{\rho}{1 + r^{nd}} \sum_{k=0}^{\infty} \frac{E_t y_{t+k}^{nd}}{(1 + r^{nd})^k}.$$

Dividing through by y_t^{nd} yields

$$1 - \frac{s_t}{y_t^{nd}} = \frac{\rho}{1 + r^{nd}} \left[1 + \sum_{k=1}^{\infty} \frac{E_t y_{t+k}^{nd} / y_t^{nd}}{(1 + r^{nd})^k} \right].$$

This equation can then be loglinearized (see the appendix of Campbell and Deaton, 1989, for details). The resulting expression is equivalent to equation (8), albeit with a slightly different constant of linearization κ (call it κ'). That is, this relationship can be re-written as

$$\frac{s_t}{y_t^{nd}} \approx - \sum_{k=1}^{\infty} \beta^k E_t \Delta \log y_{t+k}^{nd} + \kappa'.$$

We can then recover equation (26) in the text by substituting in our new definition of s_t (equation 28) and defining $k \equiv 1 + \kappa'$.

II. Data Sources and Definitions

This portion of the appendix gives the sources and definitions of the series used in the paper's empirical analysis. Note that all consumption, wealth, and income variables are expressed in per-capita terms using the population measure described below.

Consumption expenditures: Total personal consumption expenditure is taken from the National Income and Product Accounts (NIPAs). Consumption of nondurables and services is computed by combining NIPA personal consumption expenditures on nondurable goods with NIPA personal consumption expenditures on services. Real measures are combined using a Fisher chain-aggregation formula that replicates the procedure used by the Bureau of Economic Analysis in producing the National Income and Product Accounts.

Consumption prices: Price indexes are defined as implicit deflators (*i.e.*, as ratios of nominal series to corresponding real series).

Wealth: All data are taken from the Flow of Funds Accounts of the Board of Governors of the Federal Reserve System, Table B.100. Wealth is defined as household net worth less stocks of consumer durable goods. Flow of Funds wealth measures are expressed on an end-of-period basis; we therefore associate the $t - 1$ value of the data with period t wealth (that is, with A_t) in order to obtain a start-of-period measure.

Disposable income: NIPA disposable income (Table 2.1, line 25). Following Blinder and Deaton (1985), we reduce disposable income in 1975:Q2 by \$32.5 billion (at an annual rate) in order to remove the effect of the 1975 tax rebate from measured income.

Labor income: We define labor income as wage and salary disbursements (NIPA Table 2.1, line 2) *plus* transfers to persons (line 16) *plus* other labor income (line 9) *minus* personal contributions for social insurance (line 23) *minus* labor taxes. Labor taxes are defined by imputing a share of personal tax and nontax payments (line 24) to labor income, with the share calculated as the ratio of wage and salary disbursements to the sum of wage and salary disbursements, proprietors' income (line 10), and rental (line 13), dividend (line 14), and interest (line 15) income. Note that personal tax and nontax payments are adjusted for the effect of the 1975 tax rebate.⁹

Population: Population from NIPA Table 8.7, line 16. (Note that this is the population measure used by the Bureau of Economic Analysis to compute official per-capita income and consumption data.)

⁹The definition of labor income used in this paper differs from the measure used by Blinder and Deaton (1985), which was in turn employed by Campbell (1987) and Campbell and Deaton (1989). We re-did the analysis of Section 5 using two variants of an updated Blinder-Deaton income series, and found that the substantive conclusions were unaffected by using these alternative measures.

Table 1: Results from Rainy-Day Saving Equations

		Model regressors		
		$\Delta y_{t-i,i=1}$ $\left(\frac{s_{t-i}}{y_{t-i}}\right)_{i=1}$	$\Delta y_{t-i,i=1,\dots,4}$ $\left(\frac{s_{t-i}}{y_{t-i}}\right)_{i=1}$	$\Delta y_{t-i,i=1,\dots,4}$ $\left(\frac{s_{t-i}}{y_{t-i}}\right)_{i=1,\dots,4}$
<i>A. Using Traditional Deflation</i>				
<i>Labor income:</i>	Coefficient sum	0.195** (0.070)	0.306** (0.115)	0.495** (0.138)
	Exclusion restriction	0.006	0.044	0.001
<i>Saving rate:</i>	Coefficient sum	-0.037* (0.015)	-0.037* (0.015)	-0.031* (0.015)
	Exclusion restriction	0.013	0.013	0.003
<i>B. Using Our Deflation Procedure</i>				
<i>Labor income:</i>	Coefficient sum	0.210** (0.072)	0.342** (0.121)	0.521** (0.138)
	Exclusion restriction	0.004	0.024	0.000
<i>Saving rate:</i>	Coefficient sum	-0.016 (0.020)	-0.023 (0.020)	-0.004 (0.021)
	Exclusion restriction	0.417	0.265	0.013

Note: Standard errors in parentheses; **/*/^a denotes significant at 1/5/10 percent level, respectively. “Exclusion restriction” gives the p -value from an F -test that lags of given variable can be excluded from the model. Regressions include indicated number of lags of saving rate and log-differenced labor income; dependent variable in all equations is log-differenced labor income. See text for additional details.

Table 2: Rainy-Day Saving Equations, Regression-Based Capital Income

		Model regressors		
		$\frac{\Delta y_{t-i,i=1}}{\left(\frac{s_{t-i}}{y_{t-i}}\right)_{i=1}}$	$\frac{\Delta y_{t-i,i=1,\dots,4}}{\left(\frac{s_{t-i}}{y_{t-i}}\right)_{i=1}}$	$\frac{\Delta y_{t-i,i=1,\dots,4}}{\left(\frac{s_{t-i}}{y_{t-i}}\right)_{i=1,\dots,4}}$
<i>Labor income:</i>	Coefficient sum	0.171* (0.075)	0.247 ^a (0.135)	0.252 ^a (0.145)
	Exclusion restriction	0.024	0.156	0.278
<i>Saving rate:</i>	Coefficient sum	0.028 (0.023)	0.020 (0.026)	0.019 (0.026)
	Exclusion restriction	0.224	0.433	0.880

Note: Standard errors in parentheses; **/*/^a denotes significant at 1/5/10 percent level, respectively. “Exclusion restriction” gives the p -value from an F -test that lags of given variable can be excluded from the model. Regressions include indicated number of lags of saving rate and log-differenced labor income; dependent variable in all equations is log-differenced labor income. See text for additional details.

Figure 1
Ratio of Consumption of Nondurables and Services to Total Consumption




Figure 2
Ratio of Consumption to Labor Income


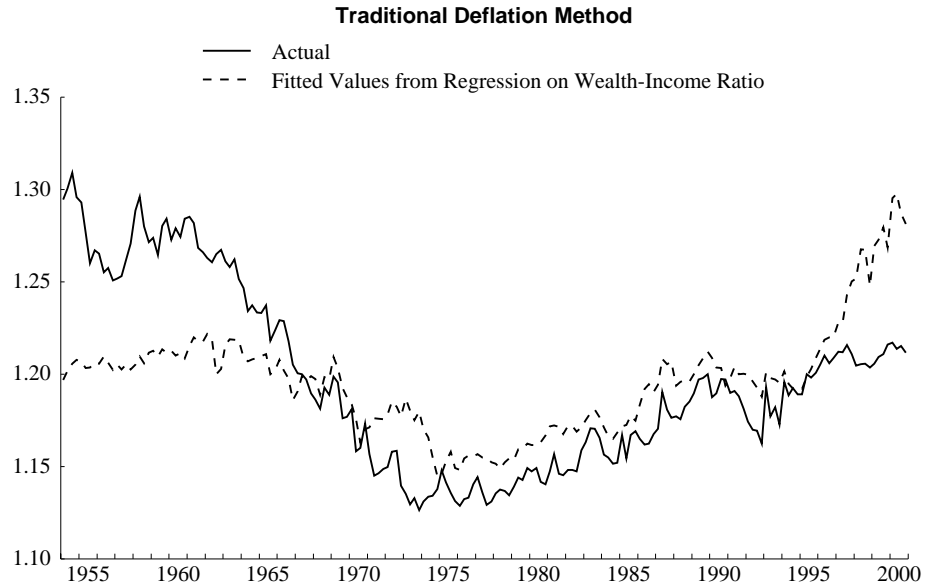


Figure 3
Two Measures of the Saving Rate




Figure 4
Regression-Based Capital Income Measurement and the Saving Rate

